



US009172556B2

(12) **United States Patent**
Del Signore et al.

(10) **Patent No.:** **US 9,172,556 B2**
(45) **Date of Patent:** ***Oct. 27, 2015**

(54) **METHOD AND APPARATUS FOR ROUTING BETWEEN FIBRE CHANNEL FABRICS**

USPC 370/351-430
See application file for complete search history.

(75) Inventors: **Christopher A. Del Signore**, San Jose, CA (US); **Vineet M. Abraham**, San Jose, CA (US); **Satish K. Gnanasekaran**, San Jose, CA (US); **Pranab Patnaik**, Sunnyvale, CA (US); **Vincent W. Guan**, Saratoga, CA (US); **Balakumar N. Kaushik**, Sunnyvale, CA (US)

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,363,367	A	11/1994	Kobayashi et al.
5,400,325	A	3/1995	Chatwani et al.
6,269,404	B1	7/2001	Hart et al.
6,285,276	B1	9/2001	Nedele et al.
6,339,842	B1	1/2002	Fernandez et al.
6,401,128	B1	6/2002	Stai et al.
6,470,007	B1	10/2002	Berman
6,529,963	B1 *	3/2003	Fredin et al. 710/1

(Continued)

(73) Assignee: **Brocade Communications Systems, Inc.**, San Jose, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 282 days.

This patent is subject to a terminal disclaimer.

OTHER PUBLICATIONS

"Fibre Channel Methodologies for Interconnects (FC-MI) Rev 1.8," NCITS working draft proposed Technical Report; Sep. 28, 2001.

(Continued)

(21) Appl. No.: **13/284,778**

Primary Examiner — Fan Ng

(22) Filed: **Oct. 28, 2011**

(74) *Attorney, Agent, or Firm* — Blank Rome, LLP

(65) **Prior Publication Data**

US 2012/0044933 A1 Feb. 23, 2012

(57)

ABSTRACT

An interfabric link between two separate Fiber Channel fabrics so that devices in one fabric can communicate with devices in another fabric without requiring the merger of the two fabrics. The interfabric switch performs a conversion or a translation of device addresses in each fabric so that they are accessible to the other fabric. In a first embodiment the external ports of the interfabric switch are configured as E_ports. A series of internal ports in each interfabric switch are joined together forming a series of virtual or logical switches. In a second embodiment the external ports are configured as NL_ports and the connections between the virtual switches are E_ports. The virtual switches in the interfabric switch match domains with their external counterparts so that the virtual switches effectively form their own fabric.

Related U.S. Application Data


(63) Continuation of application No. 10/356,392, filed on Jan. 31, 2003, now Pat. No. 8,081,642.

(51) **Int. Cl.**
H04L 12/433 (2006.01)

(52) **U.S. Cl.**
CPC **H04L 12/433** (2013.01)

(58) **Field of Classification Search**
CPC H04L 49/357; H04L 49/103; H04L 47/10; H04L 49/351

10 Claims, 13 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

- 6,532,212 B1 * 3/2003 Soloway et al. 370/230
 6,608,819 B1 8/2003 Mitchem et al.
 6,763,417 B2 7/2004 Paul et al.
 6,834,311 B2 * 12/2004 Rao 709/245
 6,879,593 B1 4/2005 Kunze et al.
 6,941,260 B2 9/2005 Emberty et al.
 6,985,490 B2 1/2006 Czeiger et al.
 7,068,651 B2 6/2006 Schmidt et al.
 7,103,704 B2 9/2006 Chatterjee
 7,103,711 B2 9/2006 Valdevit
 7,107,347 B1 9/2006 Cohen
 7,120,728 B2 10/2006 Krakirian et al.
 7,130,303 B2 10/2006 Hadzic
 7,206,288 B2 4/2007 Cometto et al.
 7,206,314 B2 4/2007 Liao et al.
 7,236,496 B2 6/2007 Chung et al.
 7,287,116 B2 10/2007 Iwami et al.
 7,292,567 B2 11/2007 Terrell et al.
 7,305,069 B1 12/2007 Day
 7,340,167 B2 3/2008 McGlaughlin
 7,385,982 B2 6/2008 Warden et al.
 7,533,256 B2 5/2009 Walter et al.
 7,542,676 B2 6/2009 McGlaughlin
 7,577,134 B2 8/2009 Gopal Gowda et al.
 7,606,167 B1 10/2009 DeSanti et al.
 7,616,637 B1 * 11/2009 Lee et al. 370/394
 7,936,769 B2 5/2011 Chung et al.
 8,055,794 B2 * 11/2011 Shanbhag et al. 709/244
 8,081,642 B2 * 12/2011 Del Signore et al. 370/401
 8,135,858 B2 * 3/2012 Shanbhag et al. 709/244
 2002/0010790 A1 1/2002 Ellis et al.
 2002/0013848 A1 1/2002 Rene Salle
 2002/0023184 A1 2/2002 Paul
 2002/0101859 A1 8/2002 Maclean
 2002/0110125 A1 8/2002 Banks et al.
 2002/0114328 A1 8/2002 Mijamoto et al.
 2002/0116564 A1 8/2002 Paul et al.
 2002/0141424 A1 10/2002 Gasbarro et al.
 2002/0161567 A1 10/2002 Emberty et al.
 2002/0163910 A1 11/2002 Wisner et al.
 2002/0165978 A1 11/2002 Chui
 2002/0188711 A1 12/2002 Meyer et al.
 2002/0191602 A1 12/2002 Woodring et al.
 2002/0191649 A1 12/2002 Woodring
 2003/0012204 A1 1/2003 Czeiger et al.
 2003/0030866 A1 2/2003 Yoo
 2003/0037127 A1 2/2003 Shah et al.
 2003/0058853 A1 3/2003 Gorbato et al.
 2003/0061220 A1 3/2003 Ibrahim et al.
 2003/0076788 A1 4/2003 Grabauskas et al.
 2003/0084219 A1 * 5/2003 Yao et al. 710/300
 2003/0095549 A1 5/2003 Berman
 2003/0118047 A1 * 6/2003 Collette et al. 370/464
 2003/0118053 A1 6/2003 Edsall et al.
 2003/0135385 A1 * 7/2003 Karpoff 705/1
 2003/0158971 A1 8/2003 Renganarayanan et al.
 2003/0189930 A1 10/2003 Terrell et al.
 2003/0189935 A1 10/2003 Warden et al.
 2003/0189936 A1 10/2003 Terrell et al.
 2003/0210685 A1 11/2003 Foster et al.
 2004/0013092 A1 1/2004 Betker et al.
 2004/0013125 A1 1/2004 Betker et al.
 2004/0024905 A1 2/2004 Liao et al.
 2004/0024911 A1 2/2004 Chung et al.
 2004/0081186 A1 4/2004 Warren et al.
 2004/0085972 A1 5/2004 Warren et al.
 2004/0146254 A1 7/2004 Morrison
 2004/0151174 A1 * 8/2004 Del Signore et al. 370/389
 2005/0018673 A1 1/2005 Dropps et al.
 2005/0025075 A1 2/2005 Dutt et al.
 2005/0036499 A1 2/2005 Dutt et al.
 2005/0044354 A1 2/2005 Hagerman
 2005/0073956 A1 4/2005 Moores et al.
 2005/0169311 A1 8/2005 Millet et al.
 2005/0198523 A1 9/2005 Shanbhag et al.
 2005/0232285 A1 10/2005 Terrell et al.
 2006/0023707 A1 2/2006 Makishima et al.
 2006/0023708 A1 2/2006 Snively et al.
 2006/0023725 A1 2/2006 Makishima et al.
 2006/0023726 A1 2/2006 Chung et al.
 2006/0023751 A1 2/2006 Wilson et al.
 2006/0034302 A1 2/2006 Peterson
 2006/0092932 A1 5/2006 Ghosh et al.
 2006/0203725 A1 9/2006 Paul et al.
 2007/0002883 A1 1/2007 Edsall et al.
 2007/0058619 A1 3/2007 Gopal Gowda et al.
 2007/0091903 A1 4/2007 Atkinson
 2008/0028096 A1 1/2008 Henderson et al.
 2009/0290589 A1 11/2009 Green
 2012/0044934 A1 * 2/2012 Del Signore et al. 370/355

OTHER PUBLICATIONS

"Fibre Channel Methodologies for Interconnects—2 (FC-MI-2) Rev 2.60," INCITS working draft proposed Technical Report; Jun. 7, 2005

American National Standard for Information Technology; "Fibre Channel—Fabric Generic Requirements (FC-FG)" Secretariat: Information Technology Industry Council; Approved Dec. 4, 1996: American National Standards Institute, Inc.

"Fibre Channel Switch Fabric (FC-SW) Rev 3.3," NCITS working draft proposed American National Standard for Information Technology; Oct. 21, 1997.

"Fibre Channel Switch Fabric—2 (FC-SW-2) Rev 5.3," NCITS working draft proposed American National Standard for Information Technology; Jun. 26, 2001 (pp. begin to 32 & 69-94).

"Fibre Channel Fabric—3 (FC-SW-3) Rev 6.6," NCITS working draft proposed American National Standard for Information Technology; Dec. 16, 2003 (pp. begin to 30 & 91-116).

"Fibre Channel Physical and Signaling Interface (FC-PH) Rev 4.3," working draft proposed American National Standard for Information Systems; Jun. 1, 1994 (pp. begin to 32).

IP Storage Working Group, IETF, "draft-monia-ips-ifcparch-00.txt", Internet Draft Standard, Nov. 2000, pp. 1-18.

INCITS Working Draft Proposed American National Standard for Information Technology, "Fibre Channel Link Services (FC-LS) Rev. 1.5.1" Sep. 6, 2006 Secretariat: Information Technology Industry Council.

IP Storage Working Group, IETF, "draft-monia-ips-ifcp-01.txt", Internet Draft Standard, Jan. 2001, pp. 1-48.

IP Storage Working Group, IETF, "draft-chau-fcipc-ifcp-encap-00.txt", Internet Draft Standard, Feb. 2001, pp. 1-8.

IP Storage Working Group, IETF, "draft-monia-ips-ifcpenc-00.txt", Internet Draft Standard, Feb. 2001, pp. 1-7.

IP Storage Working Group, IETF, "draft-tseng-ifcpmib-00.txt", Internet Draft Standard, Aug. 2001, pp. 1-22.

IP Storage Working Group, IETF, "draft-monia-ips-ifcplcc-00.txt", Internet Draft Standard, Apr. 2002, pp. 1-44.

IP Storage Working Group, IETF, "draft-ietf-ips-ifcp-wglcc-01.txt", Internet Draft Standard, May 2002, pp. 1-45.

IP Storage Working Group, IETF, "draft-ietf-ips-ifcp-01.txt", Internet Draft Standard, Feb. 2001, pp. 1-55.

IP Storage Working Group, IETF, "draft-ietf-ips-ifcp-02.txt", Internet Draft Standard, May 2001, pp. 1-68.

IP Storage Working Group, IETF, "draft-ietf-ips-ifcp-03.txt", Internet Draft Standard, Jul. 2001, pp. 1-67.

IP Storage Working Group, IETF, "draft-ietf-ips-ifcp-04.txt", Internet Draft Standard, Aug. 2001, pp. 1-84.

IP Storage Working Group, IETF, draft-ietf-ips-ifcp-05.txt, Internet Draft Standard, Sep. 2001, pp. 1-91.

IP Storage Working Group, IETF, "draft-ietf-ips-ifcp-06.txt", Internet Draft Standard, Oct. 2001, pp. 1-91.

IP Storage Working Group, IETF, draft-ietf-ips-ifcp-07.txt, Internet Draft Standard, Nov. 2001, pp. 1-90.

IP Storage Working Group, IETF, "draft-ietf-ips-ifcp-08.txt", Internet Draft Standard, Jan. 2002, pp. 1-98.

IP Storage Working Group, IETF, "draft-ietf-ips-ifcp-09.txt", Internet Draft Standard, Jan. 2002, pp. 1-97.

(56)

References Cited

OTHER PUBLICATIONS

IP Storage Working Group, IETF, "draft-ietf-ips-ifcp-10.txt", Internet Draft Standard, Feb. 2002, pp. 1-98.
 IP Storage Working Group, IETF, "draft-ietf-ips-ifcp-11.txt", Internet Draft Standard, May 2002, pp. 1-103.
 IP Storage Working Group, IETF, "draft-ietf-ips-ifcp-12.txt", Internet Draft Standard, Jun. 2002, pp. 1-104.
 IP Storage Working Group, IETF, "draft-ietf-ips-ifcp-14.txt", Internet Draft Standard, Dec. 2002, pp. 1-104.
 NCITS Working draft proposed American National Standard for Information Technology; "Fibre Channel—Framing and Signaling (FC-FS)," Feb. 8, 2002; Secretariat: Information Technology Industry Council.
 NCITS Working draft proposed American National Standard for Information Systems; "Fibre Channel—Backbone(FC-BB-2)," Mar. 6, 2002; Secretariat: Information Technology Industry Council.
 NCITS Working draft proposed American National Standard for Information Technology; "Fibre Channel—Generic Services—4 (FC-GS-4)," Sep. 19, 2001; Secretariat: Information Technology Industry Council.
 Pendry et al.; "InfiniBand Architecture: Bridge Over Troubled Waters," Apr. 27, 2000: Illuminata, Inc.; Nashua, New Hampshire.
 IP Storage Working Group, IETF, "draft-ietf-ips-ifcp-13.txt", Internet Draft Standard, Aug. 2002, pp. 1-104.
 Weber et al.; "Fibre Channel (FC) Frame Encapsulation," Dec. 2003; The Internet Society.
 "Fibre Channel Backbone (FV-BB)," Rev 4.7; NCITS working draft proposed American National Standard for Information Systems; Jun. 8, 2000.
 "Fibre Channel Backbone (FC-BB-2)," Rev 6.0; INCITS working draft proposed American National Standard for Information Systems; Feb. 4, 2003.

Travostino et al.; "IFCP—A Protocol for Internet Fibre Channel Networking," Dec. 2002.
 Rajagopal et al.; IP and ARP over Fibre Channel; Network Working Group Request for Comments 2625; Gadzooks Networks; Jun. 1999.
 Fabric Extension Study Group, Draft Minutes, T11/04-129v0, Feb. 3, 2004.
 Pelissier et al., FR_Header Definition, T11/04-241v0, Apr. 2004.
 Frame Routing Extended Header (FR_Header), T11/05-214v0 Revision 3, Mar. 15, 2005.
 Inter-Fabric Routing Extended Header (IFR_Header), T11/05-214v1 Revision 5d, May 20, 2005.
 Pelissier et al., Inter-Fabric Routing (T11/04-520v0), Jul. 30, 2004.
 Desanti et al., Inter-Fabric Routing, T11/04-408v0, Jun. 2004.
 Transport Fabric Model Rev 1.3, T11/05-075v1, Apr. 1, 2005.
 Pelissier et al., Inter-Fabric Routing Solutions, 05-232v1, Apr. 5, 2005.
 Fabric Routing Types, T11/05-099v0, Feb. 1, 2005.
 Newman et al., Seagate Fibre Channel Interface Product manual, Mar. 21, 1997, Seagate.
 Elsbernd et al., Fabric Expansion Study Group Draft minutes 03-691-v0, Oct. 7, 2003.
 Fibre Channel Framing and Signaling-2, standard, Apr. 1, 2004, page Nos. i-11, Rev. 0.10, INCITS working draft proposed American National Standard for Information Technology.
 I2O architecture Specification, Ver. 1.5, Mar. 1997.
 "Fibre Channel Physical and Signaling Interface (FC-PH) Rev 4.3", Jun. 1, 1994, p. front cover to I-9.
 "Fibre Channel Fabric Loop Attachment (FC-FLA) Rev. 2.7", Aug. 21, 1997, p. front cover to I22.
 "Fibre Channel—Switch Fabric—2 (FC-SW-2) Rev. 5.3", Jun. 26, 2001, p. front cover to I86.
 "Fibre Channel Private Loop SCSI Direct Attach (FC-PLDA) Rev. 2.1", Sep. 22, 1997, p. front cover to I-4.

* cited by examiner




FIG. 1B
(Prior Art)




FIG. 1A
(Prior Art)

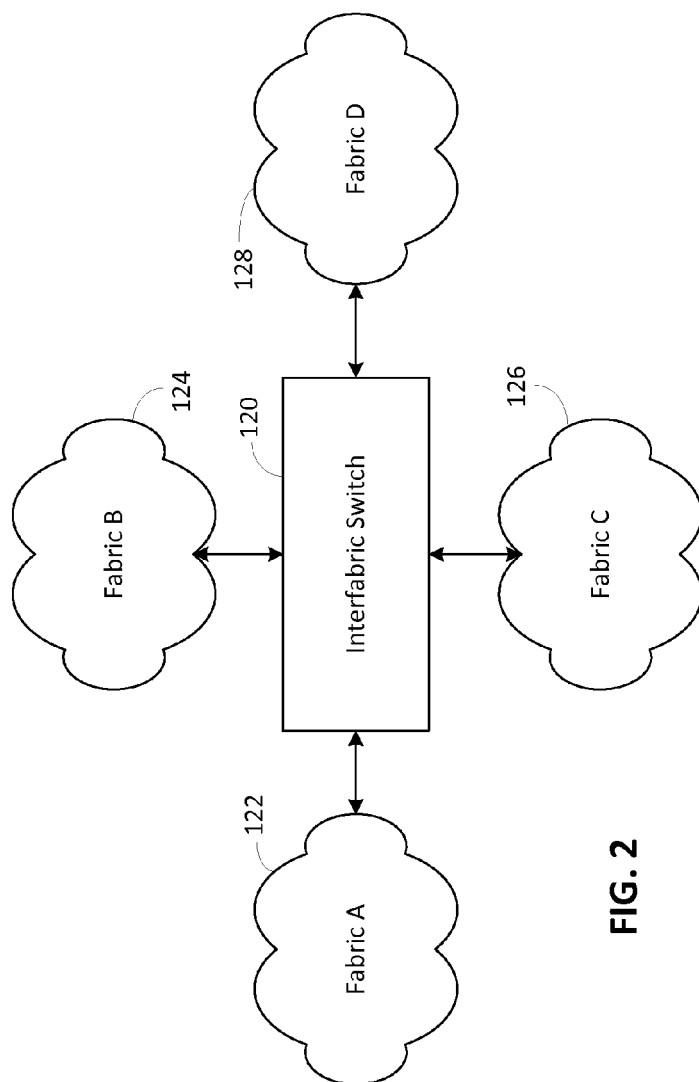


FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7





FIG. 9B

FIG. 9A






FIG. 9D





FIG. 11



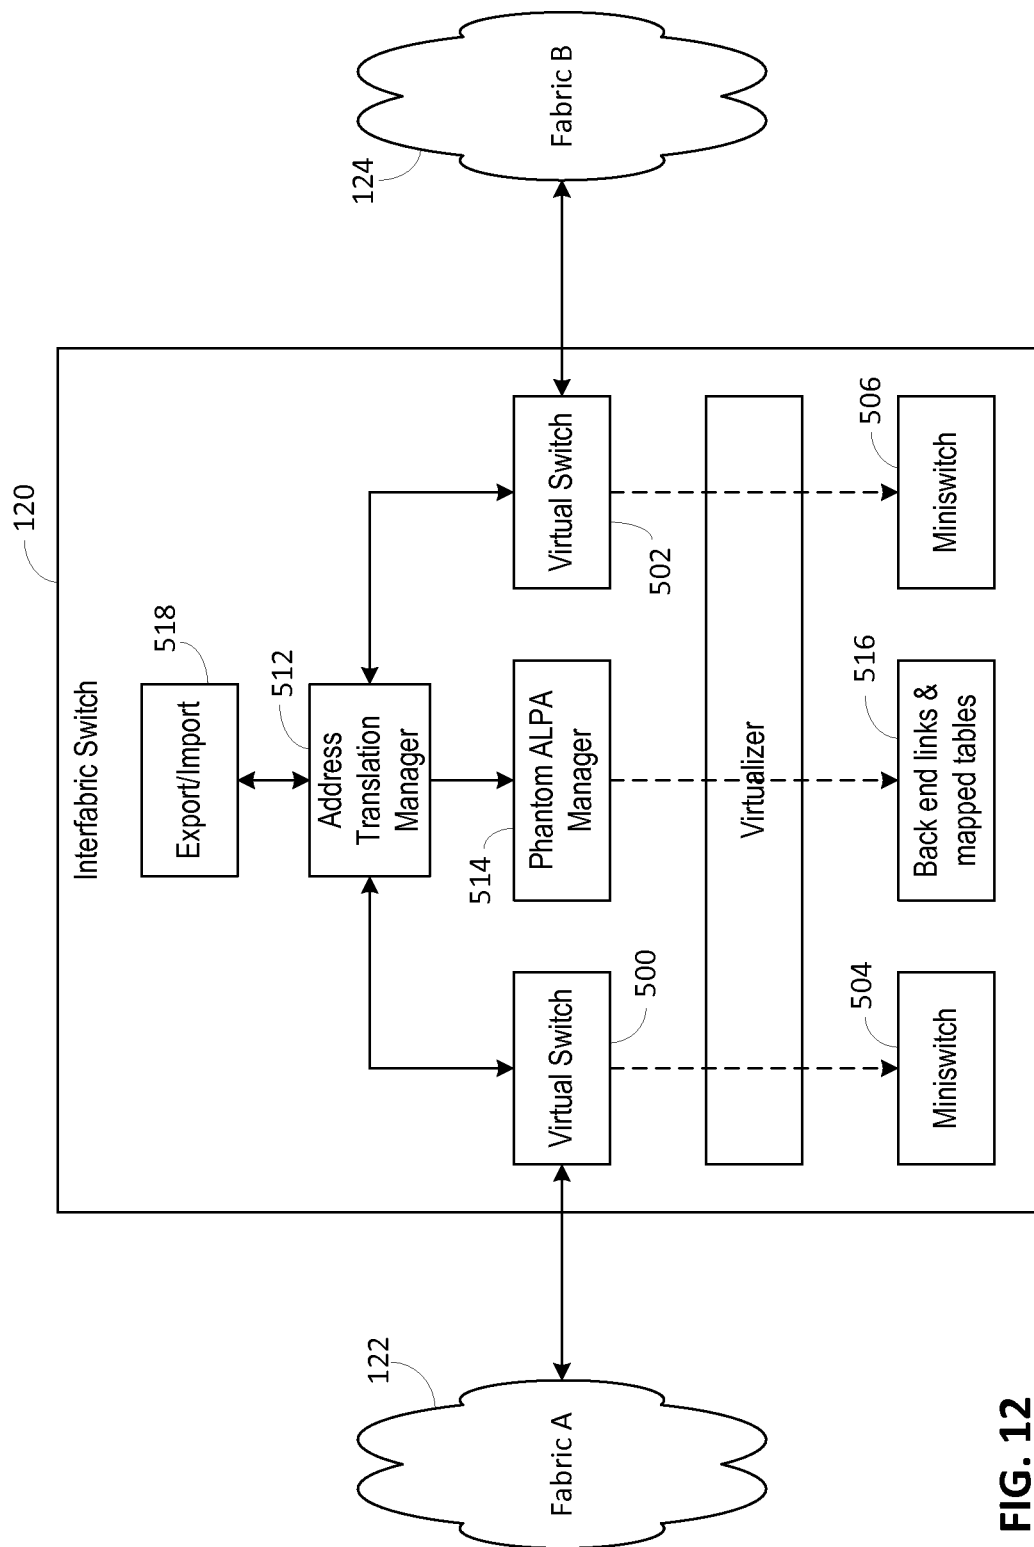


FIG. 12




FIG. 13


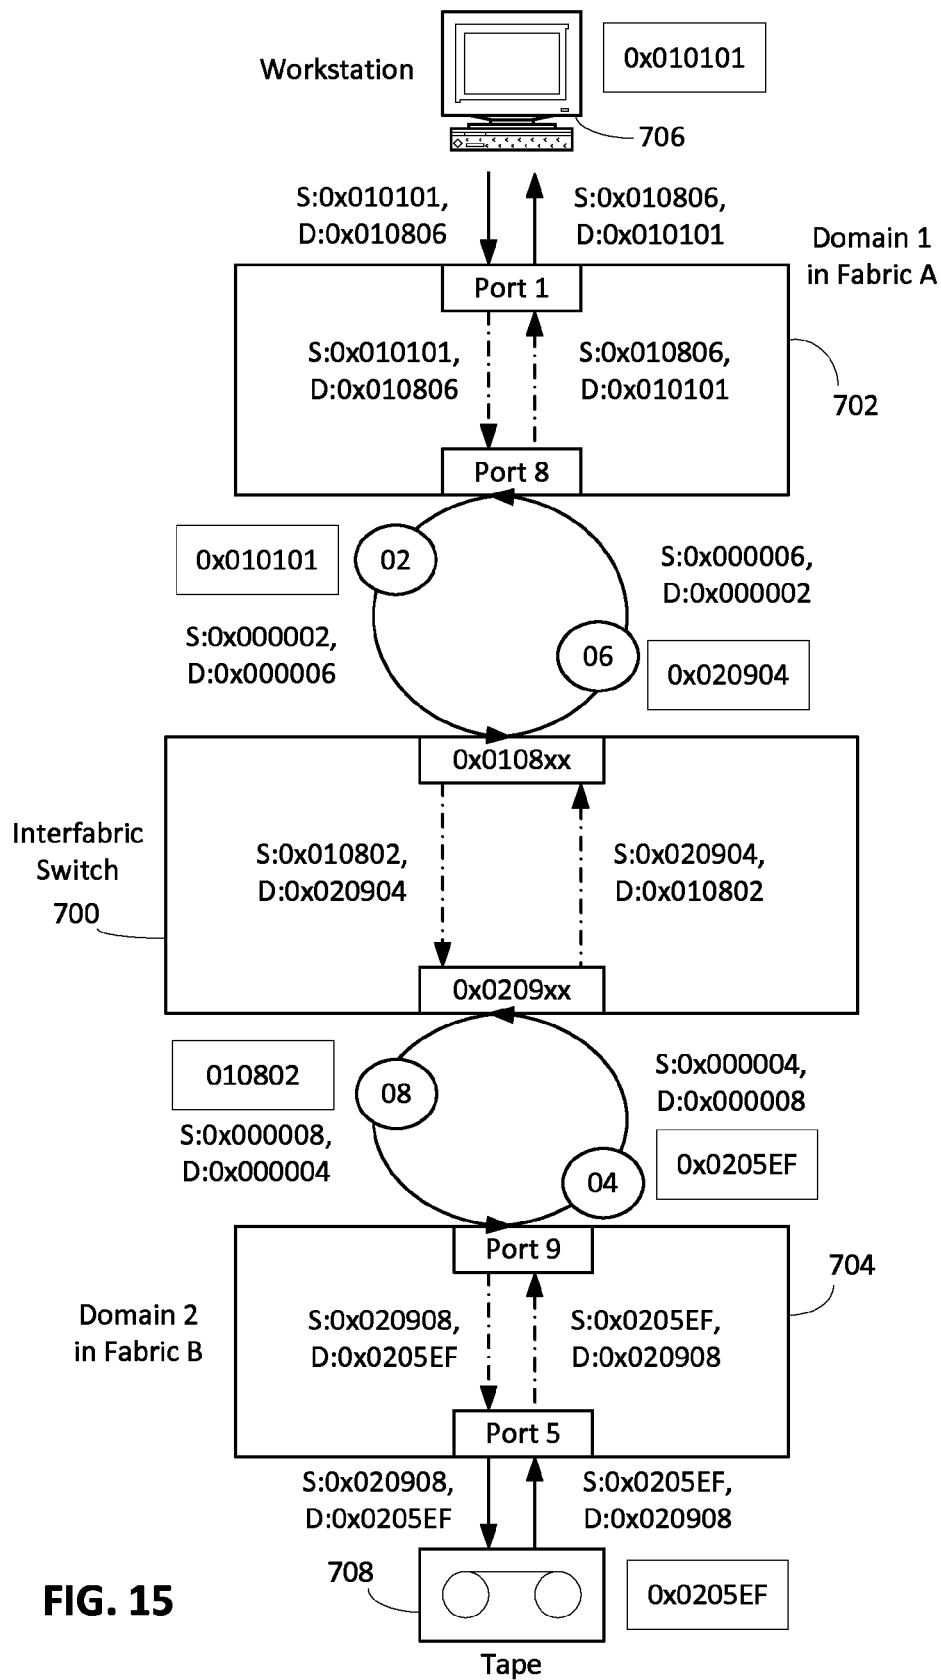


FIG. 14






FIG. 16

METHOD AND APPARATUS FOR ROUTING BETWEEN FIBRE CHANNEL FABRICS

CROSS REFERENCES TO RELATED APPLICATIONS

This application is a continuation of U.S. application Ser. No. 10/356,392 filed Jan. 31, 2003, now U.S. Pat. No. 8,081,642, entitled "Method and Apparatus for Routing Between Fibre Channel Fabrics," which is hereby incorporated by reference as though fully set forth herein.

BACKGROUND OF THE INVENTION

1. Field of the Invention

The invention relates to a method and apparatus for receiving and transmitting data in a network, and more particularly, to a method for receiving and transmitting data between separate Fibre Channel fabrics.

2. Description of the Related Art

As computing power has increased over the years, the need for high performance storage capacity has also increased. To this end, storage area networks (SANs) have been developed. Basically a SAN is an interconnection network between a series of hosts or servers and a series of storage devices. The interconnection network is very high performance to allow each of the servers to access each of the desired storage units without significant performance penalties. The use of SANs allows a more optimal use of the available storage capacity than would otherwise be the case if the storage capacity was directly attached to the particular hosts.

The preferred interconnection network for SANs is Fibre Channel. This is a high speed link system according to a series of ANSI standards. In a Fibre Channel network a series of fabrics or inter-switch connections are developed. Hosts are connected to the fabric, as are the storage units. Then the interconnected switches in the fabric provide a path or route between the host and the storage unit. Thus the development of SANs has allowed very large increases in cost effective storage capacity.

However, there are certain problems when developing networks using Fibre Channel switches. One of the problems is that there can only be 239 distinct domains in a Fibre Channel fabric. Further, there are many conditions under which the fabric will segment or break into two fabrics, so that communication between devices on the two fabrics is not possible. For example, segmentation can be caused when certain parameters associated with the particular switches are not set to the proper values. As the number of switches in the fabric grows larger, the chances of segmentation ever increase. In fact, in many cases it is not possible to maintain all of the desired switches in a single fabric. This then hinders configuration of the particular network because certain devices will not be allowed to access other devices because the two fabrics are not connected. Therefore, it is desirable to have a way to connect the two fabrics so that devices can talk across the two fabrics without requiring that the fabrics be merged or allowing the combination of the two fabrics to have a total of more than 239 domains.

BRIEF SUMMARY OF THE INVENTION

Methods and devices according to the present invention provide an interfabric link between the two separate Fibre Channel fabrics so that devices in one fabric can communicate with devices in another fabric without requiring the merger of the two fabrics. Alternatively, two fabrics with

more than a combined total of 239 domains can be created and devices can still communicate. An interfabric switch according to the present invention is connected and linked to each of the two separate fabrics. The interfabric switch then performs a conversion or a translation of device addresses in each fabric so that they are accessible to the other fabric. For example, if a host is connected to fabric A and a storage unit is connected to fabric B, the interfabric switch according to the present invention provides an effective address in fabric A for the storage unit and additionally an effective address for the host unit in fabric B. The interfabric switch then allows a link to be developed between the fabrics and transfers the data packets with the translated addresses over this link. Thus the host and storage unit can communicate as though they were in the same fabric and yet the two particular devices are in separate and distinct fabrics.

This translation is preferably done using public to private and then private to public loop address translations. Using this technique the address translation can be done at full wire speed, after initial setup, so that performance of the network is not hindered by the interfabric switch.

Two particular embodiments of the loop translation are illustrated. In a first embodiment the external ports of the interfabric switch are configured as E_ports. A series of internal ports in each interfabric switch are joined together, with the interfabric switch then having a series of virtual or logical switches. In the preferred embodiment connections from each of the two particular fabrics are provided to a different virtual switch. The internal ports forming the virtual switches are then interconnected using private loops. The use of the private loop in the internal connection is enabled by the presence of translation logic which converts fabric addresses to loop addresses and back so that loop and fabric devices can communicate. Because each port can do this translation and the private loop addressing does not include domain or area information, the change in addresses between the fabrics is simplified.

In a second embodiment the external ports are configured as NL_ports and the connections between the virtual switches are E_ports. Thus the private to public and public to private translations are done at the external ports rather than the internal ports as in the prior embodiment. The virtual switches in the interfabric switch match domains with their external counterparts so that the virtual switches effectively form their own fabric, connected to the other fabrics by the private loops.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1A is a simplified drawing to illustrate private to public mode address translation.

FIG. 1B is a diagrammatic form of the system of FIG. 1A, indicating the particular device addresses to further illustrate the translation.

FIG. 2 is a drawing of a network with an interfabric switch according to the present invention interconnecting four separate fabrics.

FIG. 3 indicates the use of two interfabric switches according to the present invention between two independent fabrics.

FIG. 4 illustrates the use of two interfabric switches according to the present invention to connect three fabrics.

FIG. 5 is a drawing illustrating the virtual devices present in an interfabric switch according to the present invention when connecting two independent fabrics.

FIG. 6 is a diagram indicating the various devices, switches and interfabric switch, including device addresses, according to the present invention to allow ease of explanation.

3

FIG. 7 is an alternative view showing a simplified version of the address translations.

FIG. 8 is a diagram indicating one possible connection of four switches in a fabric to an interfabric switch according to the present invention.

FIGS. 9A, 9B, 9C and 9D are various drawings indicating the interconnection of an interfabric switch according to the present invention with various relations of switches in a fabric.

FIG. 10 is a block diagram of an interfabric switch according to the present invention.

FIG. 11 is a drawing indicating the use of World Wide Names of various devices connected to fabrics which are connected by an interfabric switch according to the present invention.

FIG. 12 provides a block diagram of various software modules and related hardware modules in an interfabric switch according to the present invention.

FIG. 13 illustrates an alternative software module breakdown of an interfabric switch according to the present invention.

FIG. 14 illustrates various software layers present in an interfabric switch according to the present invention.

FIG. 15 is a diagram showing address flow and translation for an alternate embodiment when the interfabric switch ports are connected as NL_ports.

FIG. 16 is a diagram showing switch connections for an additional alternate embodiment.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

FIG. 1A illustrates a host device **100** connected to a switch **102**, which is in turn connected to a storage unit **104**. The host **100** includes an N_port **106** which is connected over a Fibre Channel link **108** to an F_port **110** in the switch **102**. The storage unit **104** includes an NL_port **112** which is connected by a loop **114** to an FL_port **116** in the switch **102**. Assuming that the loop **114** is a private loop so that the storage unit **104** cannot directly communicate with the host **100**, the switch **102** would include an apparatus for translating between the public and private devices in the host **100** and storage unit **104**. One preferred method for doing this is shown in U.S. Pat. No. 6,401,128 entitled "System and Method for Sending and Receiving Frames Between a Public Device and a Private Device," which is hereby incorporated by reference.

According to the patent, the switch **102** includes various tables to do public to private and private to public address conversions. The switch **102** develops a phantom loop private address for the public device and a phantom public address for the private device and maps the addresses between the public and private spaces. This is shown in more detail in FIG. 1B and in the patent. For example, in FIG. 1B the switch **102** has a domain of **01**. Assuming that port **110** is port **1** of the switch **102**, the address of port **110** is **010100** in the Fibre Channel addressing technique. Then the host **100** address is **010101**. Similarly, port **116** is configured as an FL_port. Assuming that port **116** is the port five of the switch **102**, then its address is **010500**. The storage unit **104** has an address on the loop **114** connected to the port **116** and for purposes of this explanation it is considered to have an address EF. To allow this private device storage unit **104** to be visible by the host **100** the port **116** then develops a public address for the storage unit **104** of **0105EF**. Thus the host **100** can address the storage unit **104** by using this as a destination address.

It is also required that the host **100** appears as an addressable device to the storage unit **104** on the private loop **114**.

4

This is done by having the port **116** pick an open loop port address, in this example **02**, and assign that to a phantom device representing the host **100**. Thus the storage unit **104** would address the phantom unit of the host **100** by using an address of **02** on the loop **114**. The FL_port **116** would intercept this communication addressed to an address of **02** and convert that address to **010101** indicating the full fabric address of the host mo. This address and the similarly converted storage unit **104** address of **0105EF** would be substituted in the particular packet and then the packet would be provided to the port **110** for transmission to the host **100**. Then when a communication is received from the host **100** at port **116**, the address **0105EF** is translated to the loop address EF and the address **010101** is translated to loop address **02**, so that the storage unit **104** can properly receive the frame.

FIG. 2 illustrates an interfabric switch **120** according to the present invention, interconnected between four independent fabrics **122**, **124**, **126** and **128**. Preferably the fabrics **122-128** are either fabrics that are natively understood by the interfabric switch **120** or are operating in an interoperability mode. An alternative embodiment is shown in FIG. 3 where interfabric switches **120A** and **120B** are connected between fabrics **122** and **124**. In the preferred embodiment the switches **120A** and **120B** are not operating in a redundant mode to simplify programming of the switches **120A** and **120B**. In an alternative embodiment they could be operating in a redundant manner, with a link between the switches **120A** and **120B** to allow coordinating communications.

FIG. 4 illustrates an additional embodiment with an interfabric switch **120A** interconnected between fabric **122** and fabric **124** and an interfabric switch **120B** interconnected between fabric **124** and fabric **126**. In the preferred embodiment devices in fabric C are not presented to interfabric switch **120A** but instead are presented only to devices in fabric **124** to simplify programming. However, in an alternative embodiment such extended operation can be included.

FIG. 5 is a simple illustration of operation according to the present invention. Here again, the interfabric switch **120** is connected to fabrics **122** and **124**. A host **130** is connected to fabric **122**, while a storage device **132** is connected to fabric **124**. In operation the interfabric switch **120** presents a virtual switch **134** to the fabric **122** with a virtual storage device **132'** connected to the virtual switch **134**. Thus the host **130** believes it is addressing the virtual storage device **132'**. Similarly, the interfabric switch **120** provides a virtual switch **136** to the fabric **124**, with a host **130'** connected to the virtual switch **136**. The storage device **132** thus communicates with the virtual host **130'**. The interfabric switch **120** provides a translation between the two virtual switches **134** and **136** so that communications actually go from host **130** to device **132**.

A detailed addressing example is shown in FIG. 6. The interfabric switch **120** is connected to a switch **200**, which is representative of the fabric **124**, and to a switch **202**, which is representative of the fabric **122**. In the illustrated example, port **3 208** of the interfabric switch **120** is connected to port **8 210** of the switch **202**, with both of the ports being configured in E_port mode so that the link is an interswitch link. Similarly, port **4 212** of the interfabric switch **120** is connected to a port **9 214** of switch **200**, with both of the ports being configured in an E_port mode. The workstation **206** is connected to the switch **202**, and is illustrated connecting to port **1 216**. A tape drive unit **204** is connected to port **5 218** of switch **200**. It is presumed that the switch **200** is domain **1** in fabric **122** and switch **200** is domain **2** in fabric **124**.

The interfabric switch **120**, as before, includes virtual switches **134** and **136**. In the illustrated embodiment, the virtual switch **134** is assigned domain **5** in fabric **122**, while

5

virtual switch 136 becomes domain 6 in fabric 124. Virtual switches 134 and 136 are connected by ports 220 and 222, respectively, which are configured as private loop ports so that a loop 224 results. As described above, then the workstation 206 and the tape unit 204 must have phantom addresses on the private loop 224. In the illustrated embodiment, the address 04 is provided to the tape unit 204 and the address 02 is provided for the workstation 206.

Thus the workstation 206 will address the tape drive 204 by providing a destination address of 050104 that is a full public loop address. The domain 05 indicates the virtual switch 134, the port 01 indicating the virtual port in the virtual switch 134 which in actuality is physical port 3 208. The 04 is the phantom public address of the tape 204 as provided by the private loop translation. This address of 050104 is converted by the virtual loop port 220 to a loop address of 04. This loop address of 04 in turn is translated by virtual loop port 222 to an address of 0205EF, which is the actual address of the tape unit 204 in fabric 124. This address is developed because the tape 204 is connected to port 5 218 of the switch 200, which is domain 2, and the tape unit 204 is preferably a public loop device with an actual loop address of EF. This results in an address of 0205EF for the tape unit 204. For the tape unit 204 to address the workstation 206, an address of 060102 is used. This is developed because the virtual switch 136 is in domain 6 and physical port 4 212 is virtual port 1 indicating that it is 060100 in fabric B. Then as the loop address of the workstation 206 on the virtual private loop 224 is 02, this fully presents itself to the tape unit 204 as a public address of 060102. This address of 060102 is converted by the virtual loop port 222 into a loop address of 02. Packets transmit from the virtual loop port 222 to the virtual loop port 220 are then converted from this loop address of 02 to the desired address of 010101 for the workstation. Similar flow occurs for packets from the workstation 206 to the tape unit 204.

An alternative version of this illustration of this is shown in FIG. 7. The interfabric switch 120 is connected to fabric 122 and fabric 124. A host 130 is connected to the fabric 122 and a storage device 132 is connected to fabric 124. The interfabric switch 120 includes a virtual switch 134 and a virtual switch 136 connected by a private loop 224. The virtual switch 134 is assumed to receive a domain 4 and an area or port address of 3 for the loop. The virtual switch 134 provides a phantom host 130', with a loop address of 9 while the virtual switch 136 provides a virtual storage device 132' with a loop address of 6. The virtual switch 136 is domain 3 and the loop is area or port 6. The host 130 has an address of 020600 by being on domain 2 port 6 and it would address the device 132 by using an address of 040306 to indicate domain 4, port 3 and loop device 6. The device 132 is in domain 5 and port 1 so that it has an address of 050100. Thus the virtual host 130' has an address of 030609 and this is used for addressing purposes by the device 132. The virtual switch 136 converts the 030609 address into a loop address 9 on the loop 224, which virtual switch 134 then translates to 020600 indicating the host 130.

The drawings and explanations of FIGS. 6 and 7 have used a private loop as the means for connecting the virtual switches. This was done to simplify visualization. In fact, any link between the virtual switches which causes the same address translations, such as from 24 bit to 8 bit to 24 bit, can be used. In the preferred embodiment, the links are effectively normal ISLs as between E_ports, except that only private or 8 bit addressing is used on the links, with appropriate 24 bit or public address conversions at each port. Referring to U.S. Pat. No. 6,401,128, this can be done by properly programming the mapping tables and having them active in this essentially private E_port mode. The public to private translation will

6

remove the domain and area bits, which will be provided by the private to public translation in the manner described above. This is preferred because it allows the ports to be trunked to increase throughput between the virtual switches. However, if the specific components used do not provide this capability, then private loops can be used as illustrated.

FIG. 8 is the first of several illustrations of the connections of the interfabric switch 120 to particular switches inside the fabric 122. In FIG. 8 the interfabric switch 120 is connected to four individual switches 300, 302, 304, and 306 inside the fabric 122. The interfabric switch 120 is thus allowed to connect the devices, which are routed through these four switches 300-306.

FIGS. 9A, 9B, 9C, and 9D illustrate more detailed connections of the interfabric switch 120. Each of the switches in the preferred embodiment have their ports organized into quads, with the ports in the particular quads capable of being trunked as described in U.S. patent application Ser. No. 09/872,412, entitled "Link Trunking and Measuring Link Latency in Fibre Channel Fabric" by David C. Banks, Kreg A. Martin, Shunjia Yu, Jieming Zhu and Kevan K. Kwong, filed Jun. 1, 2001, which is hereby incorporated by reference. Thus interfabric switch 120 in FIG. 9A includes a quad 308 which is connected by four trunked links 310 to a quad 312 in a switch 314 in fabric 122. In the alternate embodiment shown in FIG. 9B, the interfabric switch 120 has its quad 308 connected to quads 312 and 316 in switch 314. In this case, the two links between quad 308 and quad 312 are trunked and the two links between quad 308 and 316 are trunked. In FIG. 9C, the interfabric switch 120 has its quad 308 connected to quads 312 and 316 in switch 314 and to quads 318 and 320 in switch 322. In this embodiment, none of the links can be trunked because they are in different quads. FIG. 9D then shows a more detailed breakdown of FIG. 8 where quad 308 of interfabric switch 120 is connected to quad 312 in switch 314, to quad 318 in switch 322, to quad 324 in switch 326 and to quad 328 in switch 330. Accordingly, because these are going to four separate switches none of the links can be trunked.

FIG. 10 illustrates a block diagram of an interfabric switch 120 according to the preferred embodiment. In switch 120 a processor unit 402 which includes a high performance CPU, preferably a PowerPC, and various other peripheral devices including an Ethernet module, is present. Receiver/driver circuitry 440 for a serial port is connected to the processor unit 402, as is a PHY 406 used for an Ethernet connection. A flash memory 410 is connected to the processor 402 to provide permanent memory for the operating system and other routines of the interfabric switch 120, with DRAM 408 also connected to the processor 402 to provide the main memory utilized in the interfabric switch 120. A PCI bus 412 is provided by the processor 402 and to it are connected a series of Fabric Channel miniswitches 414, 422, 428 and 434. The Fibre Channel miniswitches 414, 422, 428 and 434 are preferably developed as shown in U.S. patent application Ser. No. 10/123,996, entitled, "Fibre Channel Zoning By Device Name In Hardware," by, Ding-Long Wu, David C. Banks, and Jieming Zhu, filed on Apr. 17, 2002 which is hereby incorporated by reference. Each of the miniswitches 414, 422, 428, and 434 are thus effectively 16 port switches, with each miniswitch further broken down into four quads, as previously referenced. Four of the ports of each of the miniswitches 414, 422, 428, and 434 are connected to a series of serializers 418, 424, 430, and 436, which are then connected to media units 420, 426, 432, and 438 to present external ports for the interfabric switch 120.

To provide the necessary interconnections to represent the virtual switches in the interfabric switch 120, the

miniswitches **414**, **422**, **428**, and **434** are interconnected. Thus, four ports on miniswitch **414** are connected to four ports on miniswitch **422**, four ports on miniswitch **414** are connected to four ports on miniswitch **434** and four ports on miniswitch **414** are connected to four ports on miniswitch **428**. Similarly, four ports on miniswitch **422** are connected to four ports on miniswitch **428** and four ports on miniswitch **422** are connected to four ports on miniswitch **434**. Finally, four ports on miniswitch **428** are connected to four ports on miniswitch **434**. Thus, this provides a full direct interconnect between any of the four miniswitches **414**, **422**, **428**, and **434**. The various ports connected between the various miniswitches are configured to be private loop ports so that the miniswitches **414**, **422**, **428**, and **434** provide the private to public translations as previously described. The external ports for the interfabric switch **120** are configured as E_ports in the miniswitches **414**, **422**, **428**, and **434**. It is also noted that each of the groups of four is preferably obtained from a quad in each of the miniswitches.

Referring to FIG. **11**, this is an alternative view of the network as shown from the perspective of the configuration manager. In this view point an interfabric switch **120** is connected to fabrics **122** and **124**. Three devices **450**, **452** and **454**, each having unique worldwide names (WWNs), are connected to fabric **122**, while devices **456**, **458**, and **460**, each also having unique worldwide names, are connected to fabric **124**. This is the perspective utilized by management software to configure the mapping in the interfabric switch **120** to allow the proper address translations to occur and to provide easiest interface to system administrators. Using management software the administrator of each fabric **122** and **124** will provide the interfabric switch **120** with a list of WWNs available for export from the fabric and a list of WWNs desired to be imported into the fabric. The interfabric switch **120** then uses these export/import lists to establish the necessary phantom devices on the private loops and the translations done by the virtual switches.

FIG. **12** illustrates an alternative block diagram of the interfabric switch **120**, which includes a mixture of software and hardware modules. In the embodiment shown in FIG. **12** only two miniswitches are illustrated for simplicity. The interfabric switch **120** includes virtual switches **500** and **502**, which are connected to fabrics **122** and **124**, respectively. Virtual switch **500** effectively maps to miniswitch **504**, while virtual switch **502** effectively maps to miniswitch **506**. An address translation manager **512**, which manages the particular address translations performed in the private to public mappings in the interfabric switch **120** is connected to the virtual switches **500** and **502**. As the name servers executing on each virtual switch detect changes in connected devices in each fabric, the changes are provided to the address translation manager **512** to allow any needed changes in address translations. The address translation manager **512** is also connected to a phantom ALPA manager **514**, which maps into the various tables **516** in the miniswitches **504** and **506** which perform the actual address translations in hardware. The phantom ALPA manager **514** receives the desired translations from the address translation manager **512** and properly sets up the various tables **516**. In addition, the phantom ALPA manager **514** also receives any PLOGI events, which are trapped by filters in the system illustrated in Ser. No. 10/123, 996. The phantom ALPA manager **514** checks with the address translation manager **512** to see if they are relevant to any devices being imported from the other fabric. If so, appropriate steps are taken and the new device is incorporated. If not, normal PLOGI handling routines are invoked. The address translation manager **512** is connected to an export/

import list **518**, which is the list of devices to be exported from and imported into each fabric. The address translation manager uses the export/import list in conjunction with information from the name server and the PLOGI trapping by the phantom ALPA manager **514** to determine the actual address translations needed in the interfabric switch **120**.

A slightly different view is shown in FIG. **13**, where the interfabric switch **120** is shown with four virtual switches **550**, **552**, **554**, and **556**, which connect to respectively fabrics **558**, **560**, **562**, and **564**. Each of the virtual switches **550**, **552**, **554**, and **556**, are connected to the address translation manager module **512**, which is connected to the export/import list **518**, which in turn, is connected to an element manager **574**. The element manager **574** provides a portion of the management interface in the interfabric switch **120** and is used to load or change the export/import database **518**. The dotted lines from the various fabrics **558**, **560**, **562**, and **564** indicate read only access to the element manager **574** for the various devices in the fabrics for identification purposes.

FIG. **14** illustrates a software layer view of the interfabric switch **120**. A CPU **600** is at the lowest layer with an operating system image **602**, preferably the Linux operating system in the preferred embodiment, executing on the CPU. A fabric operating system virtualizer module **604** is operating in the operating system image **602**. The virtualizer module **604** will make necessary changes to the miniswitch drivers to allow the miniswitches to be partitioned and to other Linux components necessary to allow multiple fabric operating system instances. After the virtualizer module **604** loads and executes, four instantiations of the virtual fabric operating system (V-FOS) **608**, **610**, **612**, and **614** are loaded. One instantiation is assigned to each miniswitch in the preferred embodiment. The virtual fabric operating system is a slightly modified copy of the fabric operating system used in a conventional switch. After the four V-FOSs **608**, **610**, **612**, and **614** are executing, a phantom ALPA bridge **606** that operates on top of the fabric OS virtualizer **604** and below the V-FOSs **608**, **610**, **612**, and **614** is loaded and executed. The V-FOS **608**, **610**, **612**, and **614** instructions configure the internally connected miniswitch ports as private link ports, either FL_ports for a private loop or private E_ports for the link described above. The phantom ALPA bridge **606** then programs the miniswitches as necessary to develop the private connections and the necessary translations, as described above. Thus, four virtual switch instantiations are present and separated by the phantom ALPA bridge **606** so that there are in actuality four virtual switches executing on the interfabric switch **120**. Each of these virtual fabric operating system instantiations **608**, **610**, **612** and **614** thus control their respective miniswitches **616**, **620**, **622**, and **624** and perform normal switch functions, such as name server functions, except that in the preferred embodiment the virtual switches cannot act as the principal switch in a fabric.

To coordinate with FIGS. **12** and **13**, the address translation manager **512**, export/import list **518** and element manager **574** execute at the level of the fabric operating system virtualizer **604** as switch-level tasks.

In previous embodiments it has been seen that the external ports of the interfabric switch **120** are configured in an E_port mode. In an alternative embodiment as shown in FIG. **15**, the external ports of an interfabric switch **700** are configured to be NL_ports. Thus the mating or connected ports on switches **702** and **704** are configured as FL_ports, while the various other ports on the switches **702** and **704** are configured normally. For example, the port of switch **700** is connected to the illustrated workstation **706** and configured in an F_port mode,

and the port connected to the tape unit **708** in switch **704** is configured as an FL_port, as the tape unit is connected in loop mode.

The NL_ports of the interfabric switch **700** are configured as having two addresses, one public and one private. The interfabric switch **700** uses the public address to log into the connected fabric and learn the address of the connected FL_port, which it then configures as its own address. The FL_port will also detect the private address by probing as described in U.S. Pat. No. 6,353,612, which is hereby incorporated by reference. The NL_port will then create a public-private translation for the private device. The FL_port will also develop a phantom address in the connected device, which the interfabric switch **700** will determine. This is done for each fabric, so the interfabric switch **700** ends up knowing all the device public-private address translations and has addresses for the connected ports in different domains.

The interfabric switch **700** then assigns public addresses for each of the phantom devices connected to each port based on the port address. The interfabric switch **700** then effectively separates the ports into virtual switches as described above, with the domain of each virtual switch defined by the public port address. The virtual switches thus effectively form their own fabric separated from the other fabrics by the loops. The virtual switches are connected by E_ports so no address translations are necessary and the public addresses of the phantom devices are used.

In this mode the public to private translations occur between the interfabric switch **700** and the switches **702** and **704** instead of internal to the interfabric switch **700**. The address mappings are shown in detail in FIG. 15 and will not be described in detail but can be readily understood based on prior descriptions and review of the figure.

As previously mentioned with respect to FIGS. 6 and 7, private loops are used to simplify visualization but any link which causes the same public to private to public address translations can be used.

In a third variation shown in FIG. 16, two interfabric switches **802** and **804** are in fabric A and fabric B, respectively. They are connected by an interfabric link (IFL) **806**. Effectively the IFL **806** is an ISL with only private addressing, as discussed with relation to FIGS. 6, 7, and 15. Each interfabric switch **802** and **804** has a port, port 2 in both illustrated cases, connected to the link. These ports can be referred to as I_ports for simplicity, but these ports operate as E_ports except they perform the public to private and private to public address translations.

A host **130** is shown connected to port 1, an F_port, of interfabric switch **802**, which is illustrated as being domain 3. A storage device **132** is similarly connected to port 1, an F_port, of interfabric switch **804**, which is illustrated as being domain 4. Thus, the address of the host **130** is **030100** and of the storage device **132** is **040100**. The interfabric switch **804** presents the storage device **132** as a phantom storage device **132'** with a private address of 6. The interfabric switch **802** presents the host **130** as a phantom host **130'** with a private address of 5. The interfabric switch **804** translates this private address 5 to a public address of **040205**, indicating connection to domain 4, port 2, device 5. Similarly, the interfabric switch **802** translates the private address 6 as **030206** indicating domain 3, port 2, device 6. Thus addressing by the various devices occurs as in the prior examples.

The I_ports must be defined as such at switch setup or initialization for proper operation. Further, messaging must occur between each of the interfabric switches **802** and **804** to confirm that they are connected through I_ports by an IFL. Additionally, each I_port will have to keep track of all allo-

cated private addresses to prevent duplication. Ports not defined as I_ports would be initialized according to normal protocols. The interfabric switches **802** and **804** would then operate as normal switches, routing frames between ports as usual.

In this embodiment a V_FOS is not required as there are no virtual switches, but the export/import list **518**, address translation manager **512** and phantom ALPA manager **514** are still needed. This embodiment does have the possible disadvantage that it may be less clear for an administrator to use as it will be more difficult to determine which ports are the I_ports, while in the prior embodiments all the ports will perform the necessary functions.

In all of the above examples of interfabric switches, most interfabric events must be suppressed so that they do not cross between the fabrics. Basically, the only messages that are passed are RSCNs for devices which are imported into the other fabric as the devices come on line or go off line in their original fabric and various SW_ILS frames as the switches initiate operations.

Additionally, certain frames must be captured for operation by the processor on each switch. One example is a PLOGI frame so that the import and export tables can be checked and the SID or DID in the header changed if necessary. A second example are various SW_ILS frames which include SID and DID values in their payload so that the payload values can be changed. This trapping is done in normal manner, such as hardware trapping as described in Ser. No. 10/123,996.

While illustrative embodiments of the invention have been illustrated and described, it will be appreciated that various changes can be made therein without departing from the spirit and scope of the invention.

The invention claimed is:

1. A network comprising:

a first Fibre Channel fabric;

a second Fibre Channel fabric;

a first device connected to said first Fibre Channel fabric;

a second device connected to said second Fibre Channel fabric; and

an interfabric device including:

a first port connected to said first Fibre Channel fabric;

a second port connected to said second Fibre Channel fabric; and

logic coupled to said first and second ports to allow Fibre Channel packets to transfer between said first device and said second device with all switching being done using Fibre Channel addresses.

2. The network of claim 1, wherein said logic includes circuitry to perform public to private address translations and circuitry to perform private to public address translations and wherein Fibre Channel packets transferred between said first device and said second device are operated on by both said public to private circuitry and said private to public circuitry.

3. The network of claim 1, wherein said logic includes a first and a second Fibre Channel switch, said first Fibre Channel switch coupled to said first port, said second Fibre Channel switch coupled to said second port and said first and second Fibre Channel switches coupled together.

4. The network of claim 3, wherein said first and second Fibre channel switches are coupled by a private link and said public to private circuitry is contained in said first Fibre Channel switch and said private to public circuitry is contained in said second Fibre Channel switch.

5. The network of claim 4, wherein said private link is a loop.

6. The network of claim 4, wherein said first Fibre Channel switch includes private to public circuitry and said second Fibre Channel switch includes public to private circuitry.

7. The network of claim 3, wherein said first Fibre Channel switch includes a portion coupled to said first port and said portion is configured to operate as a private link and said second Fibre Channel switch includes a portion coupled to said second port and said portion is configured to operate as a private link, wherein said first Fibre Channel switch includes said private to public circuitry and said second Fibre Channel switch includes said public to private circuitry.

8. The network of claim 7, wherein said private links are loops.

9. The network of claim 8, wherein said first Fibre Channel switch includes public to private circuitry and said second Fibre Channel switch includes private to public circuitry.

10. A Fibre Channel interfabric device comprising:

a first Fibre Channel link port for coupling to a first Fibre Channel fabric;

a second Fibre Channel link port for coupling to a second Fibre Channel fabric; and

logic coupled to said first and second ports to allow Fibre Channel packets to transfer between a first device connected to the first Fibre Channel fabric and a second device connected to the second Fibre Channel fabric with all switching being done using Fibre Channel addresses.

* * * * *